

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
13 March 2003 (13.03.2003)

PCT

(10) International Publication Number
WO 03/021886 A2

(51) International Patent Classification?: H04L 12/56

(21) International Application Number: PCT/GB02/03929

(22) International Filing Date: 28 August 2002 (28.08.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0120837.0 28 August 2001 (28.08.2001) GB

(71) Applicant (for all designated States except US):
TOSHIBA RESEARCH EUROPE LIMITED [GB/GB];
32 Queens Square, Bristol BS1 4ND (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): FITTON, Michael,
Philip [GB/GB]; Toshiba Research Europe Limited, 32
Queens Square, Bristol BS1 4ND (GB). LEONG, Siew,
Chung [GB/GB]; Toshiba Research Europe Limited, 32
Queens Square, Bristol BS1 4ND (GB).

(74) Agent: MIDGLEY, Jonathan, Lee; Marks & Clerk,
57-60 Lincolns Inn Fields, London WC2A 3LS (GB).

(81) Designated States (national): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG,
SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ,
VC, VN, YU, ZA, ZM, ZW.


(84) Designated States (regional): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,
ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK,
TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished
upon receipt of that report

[Continued on next page]

(54) Title: CHANNEL ADAPTATION IN A COMMUNICATION NETWORK



(57) Abstract: The invention relates to a system having a decentralized, distributed topology for a local area network or the like. Devices in the system monitor the channel quality and if it is determined that the quality is insufficient, a re-evaluation of the channel is carried out to try to identify a better channel. This new channel is then communicated to other devices in the network so that all devices can communicate on the new channel.

WO 03/021886 A2

WO 03/021886 A2



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Channel Adaptation in a Communication Network

The present invention relates to dynamic selection of channels for communication between devices, particularly for use in wireless personal area networks such as Bluetooth.

There have been a number of short-range wireless communication systems developed; one of the most recent is known as 'Bluetooth'. This system was initially developed with the intention of replacing the cables between computers and peripheral devices, e.g. a keyboard, to allow such devices to become mobile within a short-range of the device by introducing a lower power radio link. This technology developed to encompass a large number of applications in the field of wireless communications, including laptop computers, mobile phones and other mobile devices.

The current technology allows two or more such devices to form a wireless 'piconet', a short-range communication cell, to allow them to communicate with one another. Such wireless devices currently employ a frequency hopping spread spectrum system to mitigate the effects of unknown and unpredictable interference, or 'noise'. Some systems do employ fixed frequency links that can be switched in use to an alternative frequency on instruction from a fixed access point. However, these systems are generally dictated by a single device and have crude and unreliable methods for assessing the relative quality of alternative channels. Furthermore, relying on a centralised topology such as this removes much of the flexibility of a distributed structure.

Therefore according to the present invention there is provided a communication device for communication with one or more other communication devices comprising: means for measuring the signal quality of a communication channel; comparison means for determining if the measured signal quality exceeds a predetermined value; channel

reselection means responsive to said comparison means for generating a channel reselection signal to be transmitted to said one or more other communication devices and for determining a reselection channel for further communication based upon channel quality information; and switching means for switching communication to said reselection channel and for transmitting a signal, for reception by said one or more other communication devices, indicating the reselection channel.

Each device can initiate a channel reselection whether it is currently the master or not. Since there is no one specific master device which controls channel reselection and initial selection, use can be made of the distributed topology. Thus if the channel quality is poor in one part of the network area, a device in that area can initiate and complete channel reselection even if the original channel is acceptable everywhere else in the network area.

Preferably, the device comprises update means for transmitting a signal indicating the reselection channel which is received by the other communication devices. They can then retune to the new channel as well.

In addition, because the device can initiate the reselection without reference to any other device, reselection can be carried out rapidly without delays which may occur due to all the devices forming a consensus.

The channel quality information can be obtained from a channel scanning means for scanning some or all of the available channels. It is not necessary to scan all channels. For example, only one channel in a certain band may be scanned if it is likely that interference extends across such a band. Therefore, where interference of a certain bandwidth is known to be present, channels may be scanned at intervals corresponding to the interference bandwidth.

The channel scanning means may determine the channel quality information in response to said comparison means determining that the quality is insufficient or alternatively, it may rely on information already collected prior to determining that the quality is insufficient. This clearly reduces the time overhead of scanning channels after it has

been determined to reselect the channel. Furthermore, the pre-existing information may be updated after reselection has been determined to be necessary.

The devices preferably include polling means for transmitting a polling signal to be received by one or more other communication devices. In response to such a polling signal, the devices can provide channel information back to the original device. This allows any lack of homogeneity in the network to be compensated for by getting a picture of the channel status at different locations. Again the information provided may be pre-existing information or collected in response to the polling. If pre-existing information is available at one device, this can avoid the need for the supervisor device which initiated the reselection to carry out a scan of the spectrum.

Channel information which has been collected from different devices and at different times may be less accurate than other information and so the information may be weighted to reduce the significance of less useful information, for example data which was collected some time ago.

Devices may keep a list of channels and their determined quality which is updated with new quality data periodically. This ensures that the device always has some estimate of the best channel available if reselection is to occur whether initiated by itself or another device.

There are a number of factors which can be used when determining the signal quality of the existing channel to ascertain whether the quality is below a desired level. This can include measuring at least one of: the segment error rate; the segment failure rate; the received signal strength; and the carrier to interference ratio.

Similarly there are a number of factors which can be used when determining the quality of a potential new channel to ascertain whether the channel can be used as the reselection channel. This can include measuring for each channel at least one of: the interference signal strength; the carrier to interference ratio; and the packet error rate.

The devices preferably include control means for receiving a polling signal from another communication device, producing local channel quality information and transmitting said local channel quality information for reception by said another communication device. This allows a device to act as a slave for providing information as well as a supervisor for carrying out reselection. In this way when a device is not in supervisor mode, it can still assist another device by providing local channel quality information in response to a polling signal.

The present invention provides both a device as indicated above for use in a network but also a collection of two or more devices forming a network.

Similarly, the present invention provides a method for operating a device for use with a network and a method for operating a communication network comprising a plurality of communication devices, the method comprising in each device: monitoring the quality of a communication channel at a communication device, and determining if the quality is below a predetermined level and if so: producing channel quality information for a plurality of candidate channels; generating a channel reselection signal to be transmitted to one or more of said other communication devices; determining a reselection channel for communication based upon said channel quality information; transmitting a signal, for reception by other communication devices, indicating the reselection channel; and switching communication to the reselection channel.

The present invention will now be described in more detail with reference to the attached drawings in which:

Figure 1 shows a representation of a network to which the present invention can be applied;

Figure 2 is a flow diagram showing how an embodiment of the present invention operates;

Figure 3 shows an example of measurement frequencies used during channel evaluation; and

Figure 4 is a schematic diagram of a communication device according to the present invention.

The present invention relates to a system for enhancing the communication between devices operating in a network over a communication channel. The following embodiment is described in relation to a piconet operating using a radio frequency communication link as part of a wireless personal area network (WPAN), e.g. Bluetooth.

A typical network includes a number of devices, which want to communicate with each other. Such an arrangement is shown in figure 1. These devices may be part of one or more networks with some devices connecting to more than one of these networks. For example, one network might include a laptop computer 10, a printer 11, a scanner 12 and a mobile phone 13. Another network may be formed between the mobile phone 13 and a personal digital assistant 14 and a further network may be formed between the mobile phone 13 and another PC 15.

Considering the network formed between the PC 10, printer 11, scanner 12 and mobile phone 13, when operating as in the context of a high rate Bluetooth system, it is assumed that a high rate (HR) subnet has been set up as a low rate (LR) Bluetooth piconet.

Whilst basic communication is carried using the LR piconet using a broadband system, HR communication is carried out using narrowband techniques. This allows for the higher data communication rate but is much more susceptible to narrowband fading, interference in the band and time dispersion.

The subnet 1 shown in figure 1 comprises four links A-D providing HR communications between the four devices. During normal operation data is transferred between the devices. However, if the quality of the communication channel being used deteriorates then it may be necessary to reselect an alternative channel for communication. The channel may deteriorate for a number of reasons such as the

devices have moved into a different environment, moved relative to each other, another interfering source has moved closer to the devices and so on. The result is that satisfactory communication on the current channel is no longer possible.

In order to prevent loss of communication over the channel it is necessary to monitor the channel to determine when quality has fallen below a predetermined level, so that channel reselection can be carried out. Once it has been determined that channel reselection is required then it is necessary to determine an alternative channel to use. These two steps will be described in more detail below. In addition during start up it is also necessary to determine a suitable channel for use.

When multiple devices are present, if the devices are outside the coherence distance in the environment then narrowband and wideband fading will be uncorrelated between the different links A-D. For example, the link D may be deeply faded but links A-C have good received signal power. Consequently not all the devices would determine that reselection was required. Therefore, several and possibly all of the devices in the subnet monitor the channel. If one of the devices determines that the signal quality of the channel falls below a certain level then it initiates reselection as a supervisor. Some or all of the devices may initiate reselection and if so they become the supervisor. Thus any one of such devices can begin the reselection process without having to use a decentralised system where a consensus must be reached by all the devices. This is also more convenient than having a permanent master device for controlling the channel because this restricts the operability of this type of network because it means that all devices must be able to communicate with the master. The master must always be available and cannot move away from the other devices or vice versa.

As mentioned above, this dynamic channel selection process can be applied to deal with a number of problems which may restrict signal quality, including (i) interference in the band; (ii) narrowband fading; and (iii) time dispersion. These problems are somewhat linked. A high RMS delay spread is likely to be experienced when the instantaneous narrowband channel attenuation is large. Furthermore, deep fade in the band will exacerbate problems of interference. Preferably, averaging of the measurements would be carried out to remove the impact of fast fading on the interference measurement. In

addition, because interference may not be homogenous throughout the piconet such that different interferers dominate in different parts of the piconet then several measurements can be taken for an optimal solution.

Figure 2 shows a flow chart representing the steps to be carried out for monitoring and reselecting a channel. In S1 normal communication takes place. At S2 a check is carried out to determine whether each of the devices have quality of service above a predetermined level.

The quality of service can be measured in a number of ways. However, the most favoured option is monitoring the segment failure or segment error rate whilst in the most robust mode available (BPSK mode in the case of Bluetooth). The specific segment error rate required to trigger a reselection of the channel will be based on the required quality of service for the link. Other parameters that may be used are the received signal strength (RSSI) and the carrier to interference ratio (C/I) or combinations of any of these.

The regularity at which the quality of service is monitored depends on a number of factors. The most important of these is the rate at which interferers join or leave the vicinity of the devices and the multipath fading (due to the wireless channel) which is experienced on the desired and interfering links. In most situations, the rate of change due to multipath fading (as characterised by the Doppler frequency) will dominate the birth/death process of interferers. This is not a problem where the measurement is taken and then immediately acted upon such that there is little time between the measurement being taken and being acted upon.

However, where the measurement is taken, for example during an idle period, and then stored for use subsequently, such as for reselection initiated by another device, then there is a danger that the measurements may be 'stale', i.e. out of date. This can be avoided by taking the measurements at a rate that is faster than the rate of change of the environment (dominated by the Doppler frequency as described above). This can be determined by the device itself based upon examining the difference between subsequent measurements or broadcast across the subnet by one or the other or the

devices. Alternatively less frequent measurements may be tolerated by classifying the measurements taken according to their age. Newer measurements can be allocated a higher weighting value that is decreased with time as they become increasingly out of date. The quantisation of these weighting values could be related to the number of frames/packets since the measurement was taken or otherwise. Then when the measurements are taken into account, those which are older can be arranged to have less weight than more recent measurements.

Once it is determined by one of the devices at S2 that the quality of service is below the predetermined level, then the initiating device is assigned temporary supervisor status. Similarly a device initiating a new HR link is assigned temporary supervisor status. It then sends a dynamic channel selection request (DCS_Request) to the other devices in the subnet, step S3. It is assumed that there are N devices on the subnet and that the subnet is arranged such that M devices must accept a DCS_Request before it is accepted. Preferably $M=N-1$ meaning that all other devices on the subnet must accept the DCS_Request to ensure a robust implementation. The process may be speeded up by allowing M to be less than N-1.

At S4, it is determined whether the required M acceptances (DCS_Accept) have been received. If not, a check is carried out, at S5, to determine if the maximum (N-M) number of explicit rejections (DCS_Reject) has been received. If so, normal operation is resumed and DCS does not take place. Otherwise, the process continues on to step S6 where a check to see if the maximum time (T_{DCS}) has expired. If some of the devices have not responded within this time, then the process continues without their response and they are effectively assumed to have accepted the request for DCS. If the time has not yet expired then operation reverts to step S4 and so on.

At step S7 the initiating device is established as the supervisor. As such, it is responsible for determining the new channel and for updating the other devices with the new channel information.

The simplest way of determining the new channel is for the supervisor to determine the status of the alternative channels itself and make the appropriate selection. However,

where the subnet environment is not homogenous, the supervisor may select a channel which appears to be satisfactory to the supervisor but which is not ideal for another device of the subnet. Therefore, the supervisor may additionally incorporate information from some of the other devices on the subnet or even from another source which is able to provide local channel information.

Where measurement is to take place during the reselection process, it is possible that considerable latency would be introduced. For example, to carry out a full measurement of the entire band takes the following time: -

number of channels x (switching time + measurement time)

For a Bluetooth implementation, the settling time is likely to be of the order of 200 μ s and a measurement time of 1 μ s. With a Bluetooth device performing measurements at 1MHz, the time taken for an entire sweep is around 16ms. An HR device could scan at 4MHz and so perform the same measurement in 4ms. Whilst these figures are pessimistic (times of approximately 50% may be more realistic) they do not include additional time for averaging of channel variations which may be required.

Consequently, it is preferable if the supervisor can obtain channel data from other sources as well as or even in place of its own measurements (e.g. because it is unable to make a measurement in the time available or measurement would put undue demands on the power requirements of the device). Obtaining information from other capable devices helps to ensure that the lack of homogeneity across the subnet is accounted for. The supervisor collates measurement reports from capable devices. The reports contain a list of the most suitable and/or the least suitable channels. The supervisor can then combine the gathered information to determine the suitability of each channel as a candidate for the new channel. The significance of the data from each device or within a single report can be weighted according to various factors such as accuracy or to rank the frequencies in order of suitability.

The information used may be collated in advance. If information is already available then this reduces the delay in developing the measurement information during the

dynamic channel selection process. For example, if an LR and an HR net is co-ordinated and the LR piconet reports on current channel conditions, then this information may be used. Alternatively, one or more of the devices in the HR subnet may be designated as monitoring devices and regularly rescan the entire band to maintain an up to date set of channel status data. This could be done when the device is idle. In this way when a device is idle, it can update its channel status information. The ideal situation is where every device on the piconet would have its own model of the channel, interference and environment. Therefore, the combination of the information from such devices would be able to provide a comprehensive picture of the overall parameters and status of the band.

Ideally the supervisor will have sufficient pre-measured information available to determine the new channel. However, if this is not the case then again the supervisor may utilise one or more other devices to carry out measurements for it. Such units are referred to as designated monitoring units. Such units may be in active, sniff or park modes. Ideally devices in park mode are used since the scanning operation does not affect their performance. Where they are in active mode, such scanning operation may prevent normal operation, which is undesirable. Obviously, power consumption needs to be considered, particularly in portable devices, as it is undesirable to have excessive power consumption due to scanning activities when the device is in park mode, i.e. not carrying out its normal functions.

When the decision to begin dynamic channel selection is made, the supervisor is generally in one of three positions: full knowledge of the channels (e.g. through pre-measurement); not in possession of full channel information but with sufficient time and capability to measure channels; and unable to measure channels. The selection of the algorithm for determining the channel depends upon which of these positions apply.

With full knowledge of a channel, in the case of a single link with a single measurement, the channel with lowest interference should be selected. With multiple links and hence multiple measurements, it is necessary to carry out optimisation (as described above) based upon the information from the various devices. Where the

entire band is scanned and no suitable channel can be determined then a switch to low-rate (LR) or a disconnection is initiated.

Under the second condition, where information is not available but there is sufficient time to determine the information then it is possible to conduct a sequential search across the available frequencies, i.e. sweeping the entire band. Alternatively, the channels can be scanned randomly. However, by considering the bandwidth of potential interferers, subsequent channel measurements can be performed separated by this bandwidth or greater. In this way the possibility of wasting time scanning several channels which are all adversely affected by the same interferer is avoided, hence improving the scanning rate. For example, if the likely interferer has a 20MHz bandwidth (e.g. 802.11b wireless LAN) it should be ensured that the next channel scanned is more than 20 MHz away. Additionally, where the calculation is based on a joint optimisation of carrier and interference power, then it is desirable to ensure that the gap between subsequent channels scanned exceeds the coherence bandwidth.

Also, it is generally not essential to search the entire band but simply scan until a channel that meets a minimum requirement is identified. Again if the entire band has been scanned and no suitable candidate has been identified then a drop back to LR or disconnection should be initiated.

Finally, where insufficient pre-measured information is available and where there is insufficient time to initiate a full scan then a random selection pattern can be implemented. The number of measurements that can be carried out will be determined by the latency permitted by the logical channel(s) quality of service requirements. The worst case and hence the fastest is where a random channel is selected without prior measurement. One further option is to return to the best hop frequency in the set or remain at the last one measured.

The number of times that channel reselection is initiated (and the overall channel quality) can be used to control the switch back to LR mode. In this way, if re-selection is continually taking place, for example because the entire band is poor a return to LR is initiated.

In order to determine the best channel, the channels need to be compared based on one or more parameters as a guide to their likely quality of service. Interference Signal strength indicator (ISSI) is the main factor in assessing each channel, although it does not reflect the degree of fading the desired channel(s) will experience and therefore the resulting quality. Other parameters are the carrier to interference ratio (C/I). This requires probing the band with a transmitted signal. If this is applied to all bands, it requires n^2 measurements (for n channels) as compared to n measurements for the interference only estimate. The C/I can be calculated by looking at the deviation of the known sequence from desired values or by taking separate measurements of received signal strength during reception and ISSI during guard intervals or gaps in transmission. Due to the significant processing and hence power consumption, C/I indication is preferable for a small number of devices such as a single link between two devices.

A further parameter is to measure the packet error rate for the associated LR piconet. The above parameters may be combined or used in isolation depending upon the arrangement, environment, etc.

A schematic representation of an exemplary communication device 40 of the present invention is shown in figure 4. The device comprises an antenna 41 connected to a receiver/transmitter 42 for allowing communication with other terminals. The quality of service calculator 43 monitors the quality of service of the received signal to determine if it exceeds the predetermined threshold. If the quality of service drops below this level then dynamic channel reselection must be carried out and a signal is sent to the channel quality information requestor 44 and the channel quality calculator 45. The channel quality information requestor 44 sends out signals to the other devices in the piconet to indicate that DCS is taking place and that it is temporarily assuming the role of supervisor for the piconet.

The channel quality calculator 45 begins analysing the channels to determine the local channel properties. If required, the channel quality information receiver 46 receives channel quality information from the other devices of the piconet which have been polled by the signal from the channel equality information requestor 44. The channel

quality information from the channel quality calculator 45 and/or the channel quality information receiver 46 is sent to the alternative channel selector 47 which determines the best channel for further communication based upon the received channel information and/or on stored information. The determined channel is then used to adjust the receiver/transmitter 42 for future transmission and also to send a signal to the other devices on the piconet for similar channel reassignment.

Methods for carrying out measurement of carrier and interference signal power will now be described in more detail both in terms of channel selection during setting up of a HR piconet and during channel reselection. These methods provide an example of a base level of testing, so that it can be specified that devices have at least this level of functionality. This obviously does not prevent the application of more accurate methods that perform better than those specified here.

The application of these measurements (i.e. which units have to perform them and when) is described below.

The measurement of carrier power in a slave-master link is done in Bluetooth 1 mode, in this construction. This is achieved by scanning through the band and measuring corresponding carrier power values at each defined channel following the hopping pattern of the master. The slave determines the carrier power (C) at each channel utilising the received signal strength indication (RSSI) functionality of the receiver. The measurement is done during all 1-slot packets (and first slot of multiple slot frames) from master-to-slave slots. The measured absolute received signal strengths (taking into account the possible effects of adaptive gain control) are stored on the condition that the known access code (used for synchronization and addressing in Bluetooth 1 mode) has been found, i.e. the correlation threshold level has been exceeded. These measured C values can be averaged for each 1MHz channel.

If transmitter power control is applied, it is necessary for the master to ensure sure that the transmit power of the transmitting slave is known, for example by giving powerUp commands until the transmission power (TxP) of the slave is in at its maximum (MaxTxP).

The measurement of carrier power in a slave-slave link can also be carried out to allow the path loss between slaves to be measured. This is potentially required as there is slave-to-slave communication in high-rate, which does not occur in low rate. If transmitter power control is applied, the master makes sure that the transmit power of the transmitting slave is known, for example by giving powerUp commands until the TxP of the slave is at MaxTxP.

In the next step, the low rate master commands slave A to measure slave-to-slave path loss. Next, the master transmits to slave B and slave B responds on the return slot. The slave A makes a measurement during the slave B transmission. The measurement by slave A now determines the receive power of the transmission from slave B (as well as all inter piconet interference). It may be necessary to repeat the measurement a number of times to ensure sufficient averaging. Finally, slave A reports its measurements (receive power levels) to the master, who already knows the transmit power of the slave B, hence the path loss can be estimated.

Determination of the interference level in the frequency band is one of the most important measures used in DCS, and therefore in the selection of the best channel on which to position the HR subnet. This measurement can be performed in low-rate mode for channel selection. Two methods for performing the interference plus noise are described here:

The interference power is measured during slave-to-master slots. Again, it might be beneficial to average multiple measurements per channel. In order to avoid measuring the slave-to-master transmission itself or its spectral leakages, an appropriate frequency offset between the slave-to-master frequency channel and the frequency channel to be measured needs to be used. This frequency offset value has to be high enough so that the transmitted power leaking to the adjacent channel does not affect the measurement results (see figure 3).

The interference power is measured during the guard period between transmissions. Either the retune to the new frequency is delayed, and the receiver measures the

interference (using RSSI) at the current frequency, or the receiver is retuned to the new frequency and the measurement is performed there. This approach has the advantage that it is unlikely to require additional hardware as simultaneous operation and measurement are not required.

There is currently no specification of the transmit power mask to use (only the period when data is transmitted is specified). Therefore, a suitable interval is inserted to ensure that it is the interference power that is observed rather than the desired transmitter ramping its power. Consequently, this imposes even more of a restriction on the settling time of the hopping synthesizer. The guard period then encompasses

$$\text{guard_period} > \text{synth_settling} + \text{interference_measurement} + \text{ramping_interval}$$

It should be noted that this interval will be further reduced by the introduction of certain radio improvements features, such as the extended header.

Referring back to figure 2, at step S13 the supervisor determines the new frequency channel. This may be with information obtained from other devices by following the optional steps S8-S12. In these steps, the supervisor issues a request to other devices. Capable devices return their reports on the channel's status. When all the reports have been returned or after a predetermined time, $T_{\text{DCS_measurement}}$, the process continues to step S13 where the supervisor determines the new channel. The supervisor broadcasts the new channel to the other devices. If an acknowledgement is required then it is necessary to wait for all the devices to acknowledge the change of channel (S15-S17). If all the devices have not acknowledged the change of channel within a predetermined period of time, $T_{\text{DCS_ACK}}$, then the supervisor resends the change of channel command. This is repeated a number of times (S19) after which the change of channel is aborted and the system returns to 'Normal operation'. Once all the devices have acknowledged, all devices switch to the new channel.

It is preferable to avoid waiting for acknowledgement from all devices, as it is possible this may delay or inhibit changing of channel if one of any of the devices fails to respond. Instead, if no acknowledgement is required, those devices which receive the

channel change information make the switch and those that don't are no longer able to communicate and so drop down to the LR mode. The LR master is updated with the new HR channel so that new devices can be directed to the right frequency and so that devices which have lost track and have dropped down to LR mode can return to HR mode.

There are two main scenarios when it is necessary to perform a channel selection operation. These are during start up of a high rate subnet (channel selection) and during high rate transmission (channel reselection). A device initiating the high rate subnet would perform initial channel selection but according to the present invention other of the devices of the piconet may be able to carry out channel reselection.

The dynamic channel selection process will now be described in more detail. In Bluetooth, the high-rate RF channel is based on 4 MHz wide channels in the 2.45 GHz ISM band. Before a high-rate channel can be established, a Bluetooth FH piconet must exist. Units in the FH piconet can negotiate for the high-rate channel. The location of the high-rate RF channel in the spectrum is adaptive and responsive to occupation by other ISM users.

For the high-rate mode, 77 overlapping channels are defined. The carrier spacing is chosen to be 1 MHz to obtain the highest resolution for positioning the RF channel in the ISM spectrum. Instead of hopping, dynamic channel selection (DCS) is used. Before channel establishment, the participating units carry out measurements in the entire ISM band. Based on these measurements, the most appropriate part of the spectrum is selected to position the high-rate channel. Sufficient averaging can be applied to cancel short-term interference (e.g. other devices carrying out frequency hopping cannot be avoided but their influence can be averaged out).

There are a number of different measurement scenarios. In each case, an ordered list of channel quality can be formed, describing the most suitable carriers that can be used for the high rate link. The first carrier on the list is the one believed to give the best transmission result, the second one is the carrier believed to give the second best result

and so on. The ordered list is preferably distributed before the devices enter the HR mode, and the best carrier is used as the initial channel of the HR subnet.

In this system, there are two main ways to measure the channel:

- measure carrier (C) and interference (I) + noise(N) levels
- measure interference + noise level

Furthermore, it is possible to perform these measurements at a single point (e.g. done by the current master) or at different points using multiple devices that intend joining the high rate piconet. Using multiple measurements is more accurate as it will reflect the non-homogeneous nature of interference in the prospective high rate subnet. However, it is more complex as described below.

In all cases, the first part of the link adaptation consists of using all available measurement information to create an ordered list of the 77 possible carriers that can be used for the high rate link. As described below, this ordered list can be used in dynamic channel reselection.

Regardless of the measurement method and number of measurements that are reported, the ordered channel quality list is formed. The low rate master that is setting up the HR subnet (or a proxy) collates all the measurement reports. When measurement reports are exchanged between devices, each device scores all 77 channels with a rating of zero (the best) to M (the worst) where $M=3$ for example. The overall ordered list is then formed by summing together all available measurement reports. At this stage it is possible to scale each measurement report if required, to indicate the quality of the estimate.

Preferably, the minimum baseline requirement for performing initial channel selection when setting up a high rate subnet is based upon interference power measurements. Where the test strategy is designed assuming that this method is employed, if a more primitive (and insufficient) technique is used to select the new HR channel, it will fail

the specified tests. The interference signal strength measurement can be performed by a proxy depending on the implementation.

The ordered channel quality list is formed from a set of measurements performed at a single point in the prospective HR subnet. For example, if the quality indication for each channel is 0 to 3, 0 corresponds to negligible observed interference, 3 to very high interference levels. Suitable averaging can be applied. The purpose of this approach is to ensure that the HR subnet is placed on a channel which has sufficiently low interference contribution.

A slightly extended version of the interference only measurement system described previously uses carrier and interference power measurements for initial channel selection. Here it is assumed that the measurement point (the low rate master or proxy) will be involved in the high rate subnet that is being initiated. This device performs carrier power measurements on one or more links to other devices that will also be participating in the HR subnet.

In the case of the low rate master performing the measurement (and participating in the HR subnet) this corresponds to measuring of the carrier on appropriate slave-master links. Here, the quality indicator for each channel will be based on the ratio between carrier and interferer,

The advantage of this approach is that it reflects to a certain extent the narrowband fading that will occur in the subnet. Obviously, if more than two devices are participating in the HR subnet, a single measurement point will not reflect the channel quality which will occur on all links. For example if device 10 in figure 1 is monitoring the carrier power received from devices 11 and 13, links A and B are characterized, but not link D. Consequently this approach is most suitable if the intended HR subnet will initially be set-up with just a single physical link between two devices.

A further method may be implemented which utilises interference measurements from multiple participating HR devices. This method more accurately reflects that devices in different locations within the subnet will experience different interference levels. When

a request to initiate a high rate subnet is issued, all devices intending to join the new subnet have the option to send an interference measurement to the low rate master. As described previously, a standardized "scoring" scheme for exchanging this data is used, for example 0 to 3 with 0 being the highest quality and 3 being the worst.

In the simplest case this could correspond to participating devices submitting an interference measure to the device collating measurement reports. This method would at least reflect the potentially non-homogeneous nature of interference.

A more accurate method is for relevant devices to report carrier-to-interference ratios. However, there is an obvious increase in complexity and monitoring required for these measurement reports.

As described above, an ordered list of the suitability of all available channels is formed during initial channel selection. This system can then be employed during dynamic channel reselection if the quality of the current channel degrades.

Preferably, the ordered channel quality list is kept as up to date as possible in order to serve its purpose, as described below. This is achieved using one of a number of methods. For example one or more units on the BT2 link also are active on the BT1 link, so that measurements can be performed for the entire band. If the ordered channel quality list is updated, then the previous values can be incorporated into the new calculation as described above, using a suitable weighting to provide some memory of channel quality. The quality for the used carrier is of course monitored on a packet by packet basis. If the ordered channel quality list is updated, this information should be communicated over the HR link (on a network configuration interval, see [bt-bb]).

Ideally, this means that the units on the HR link know which carrier to change to. The supervisor will broadcast a message indicating when a change in channel will occur (and to which channel). However, it is possible that the broadcast message may not be received by all devices. Therefore, the units need to know when the new carrier will be used. The network synchronisation slot or a priority slot on the HR link serve this purpose. If no network activity is observed during a certain predetermined time (e.g. a

network synchronisation slot is missed), both the transmitter and the receiver change to the next carrier at the next priority slot. If the change of carrier is successful, the fact that a change of carrier has taken place is then communicated to the low rate master (if the master is not already aware of this).

Naturally, a change of carrier might not give the desired result, since the new carrier might also be bad. If this is the case, one will have to fall back to the low rate (hopping) mode and simply set up a new high rate link in the same way as the previous one was established.

Dynamic channel selection should be initiated based upon the segment error rate (SER) of a particular packet or sequence of packets failing predetermined quality of service parameters. The segment error rate that causes dynamic channel selection can be varied for different links, reflecting different quality requirements. Furthermore, the level may be associated with a particular modulation scheme.

For example, an application may set the required SER to be 10% with QPSK, as the throughput in BPSK mode is insufficient. Consequently, if the SER is poorer than this level, dynamic channel reselection is initiated rather than link adaption to DBPSK mode.

The present invention has been described above in relation to a WPAN such as Bluetooth. However, this invention can be applied to any kind of network where a distributed topology is used, to allow decentralised control for allocation and switching of a communication channel. This allows devices to move freely in and out of a network coverage area as no device is necessary to the operation, i.e. as a master, of the network. Similarly, any device can join a network and reselect the channel if the current channel is not providing sufficient quality. Other advantages include the ability to obtain a picture of the channel environment across the network area when selecting a new channel rather than simply selecting a channel which is of good quality in the vicinity of the master or supervisor.

CLAIMS:

1. A communication device for communication with one or more other communication devices comprising:
means for measuring the signal quality of a communication channel;
comparison means for determining if the measured signal quality exceeds a predetermined value;
channel reselection means responsive to said comparison means for generating a channel reselection signal to be transmitted to said one or more other communication devices and for determining a reselection channel for further communication based upon channel quality information; and
switching means for switching communication to said reselection channel and for transmitting a signal, for reception by said one or more other communication devices, indicating the reselection channel.
2. A communication device according to claim 1, wherein at least some of said channel quality information is obtained from a channel scanning means.
3. A communication device according to claim 2, wherein said channel scanning means determines said channel quality information in response to said comparison means.
4. A communication device according to any one of the preceding claims, further comprising polling means for transmitting a polling signal for reception by one or more other communication devices and for receiving external channel quality information from at least one of said one or more other communication devices to provide at least some of said channel quality information.
5. A communication device according to any one of the preceding claims, wherein said channel reselection means stores a list of channel quality data which is updated using the channel quality information and said reselection channel is determined based upon said channel quality data.

6. A communication device according to any one of the preceding claims, wherein the means for measuring the signal quality measures at least one of: the segment error rate; the segment failure rate; the received signal strength; and the carrier to interference ratio.

7. A communication device according to any one of the preceding claims, wherein the channel reselection means determines a reselection channel based upon one or more of: the interference signal strength; the carrier to interference ratio; and the packet error rate, of each channel.

8. A communication device according to any one of the preceding claims, further comprising:

control means for receiving a channel reselection signal from another communication device, producing local channel quality information and transmitting said local channel quality information for reception by said another communication device.

9. A communication device according to claim 8 wherein said control means produces said local channel quality information in response to said channel reselection signal.

10. A communication device according to claim 8 wherein said control means produces said local channel quality information based on previously stored local channel quality information.

11. A communication device for communication with one or more other communication devices comprising:

control means for receiving a channel reselection signal from another communication device, producing local channel quality information and transmitting said local channel quality information for reception by said another communication device.

12. A communication device according to any one of claims 8 to 11 further comprising switching means for receiving reselection channel information and switching the current communication channel to the reselection channel.
13. A communication network comprising a plurality of communication devices according to any one of the preceding claims.
14. A communication network comprising a plurality of communication devices, each device comprising:
means for measuring the signal quality of a communication channel; -
comparison means for determining if the measured signal quality exceeds a predetermined value;
channel reselection means responsive to said comparison means for determining a reselection channel for further communication based upon channel quality information;
means for transmitting information on said reselection channel for reception by the other devices; and
switching means for switching communication to said reselection channel.
15. A communication network according to claim 14 wherein said switching means is responsive to reselection channel information, transmitted by the transmitting means of another device, to switch to communicating on the reselection channel.
16. A method for operating a communication network comprising a plurality of communication devices, the method comprising in each device:
monitoring the quality of a communication channel at a communication device, and
determining if the quality is below a predetermined level and if so:
generating a channel reselection signal to be transmitted to one or more of the other communication devices;
producing channel quality information for a plurality of candidate channels;
determining a reselection channel for communication based upon said channel quality information;
transmitting a signal, for reception by other communication devices, indicating the reselection channel; and

switching communication to the reselection channel.

17. A method according to claim 16, further comprising scanning candidate channels to provide said channel quality information.

18. A method according to claim 16 or 17, wherein said scanning is carried out in response to the determination that the quality is below a predetermined level.

19. A method according to any one of claims 16 to 18, further comprising, in response to the determination that the quality is below a predetermined level:

obtaining external channel quality information from at least one other of said communication devices in response to said channel reselection signal to provide at least some of said channel quality information.

20. A method according to claim 19, further comprising storing a list of channel quality data which is updated using the channel quality information and determining said reselection channel based upon said list of channel quality data.

21. A method according to any one of claims 16 to 20, wherein the monitoring of the quality of a communication channel includes measuring at least one of: the segment error rate; the segment failure rate; the received signal strength; and the carrier to interference ratio.

22. A method according to any one of claims 16 to 21, wherein determining of said reselection channel includes measuring at least one of: interference signal strength; the carrier to interference ratio; and the packet error rate, of each channel.

23. A method according to any one of claims 16 to 22, further comprising producing localised channel quality information in a communication device in response to receiving a channel reselection signal from another communication device, and transmitting said local channel quality information for reception by said another communication device.

25

24. A method according to claim 23 wherein said local channel quality information is produced after said channel reselection signal is received.
25. A method according to claim 23 wherein said local channel quality information is produced based on previously stored local channel quality information.
26. A communication device substantially as described herein with reference to the attached drawings.
27. A method of controlling a communication network substantially as described herein with reference to the attached drawings.
28. A communication system substantially as described herein with reference to the attached drawings.

1/5




Fig. 1

2/5

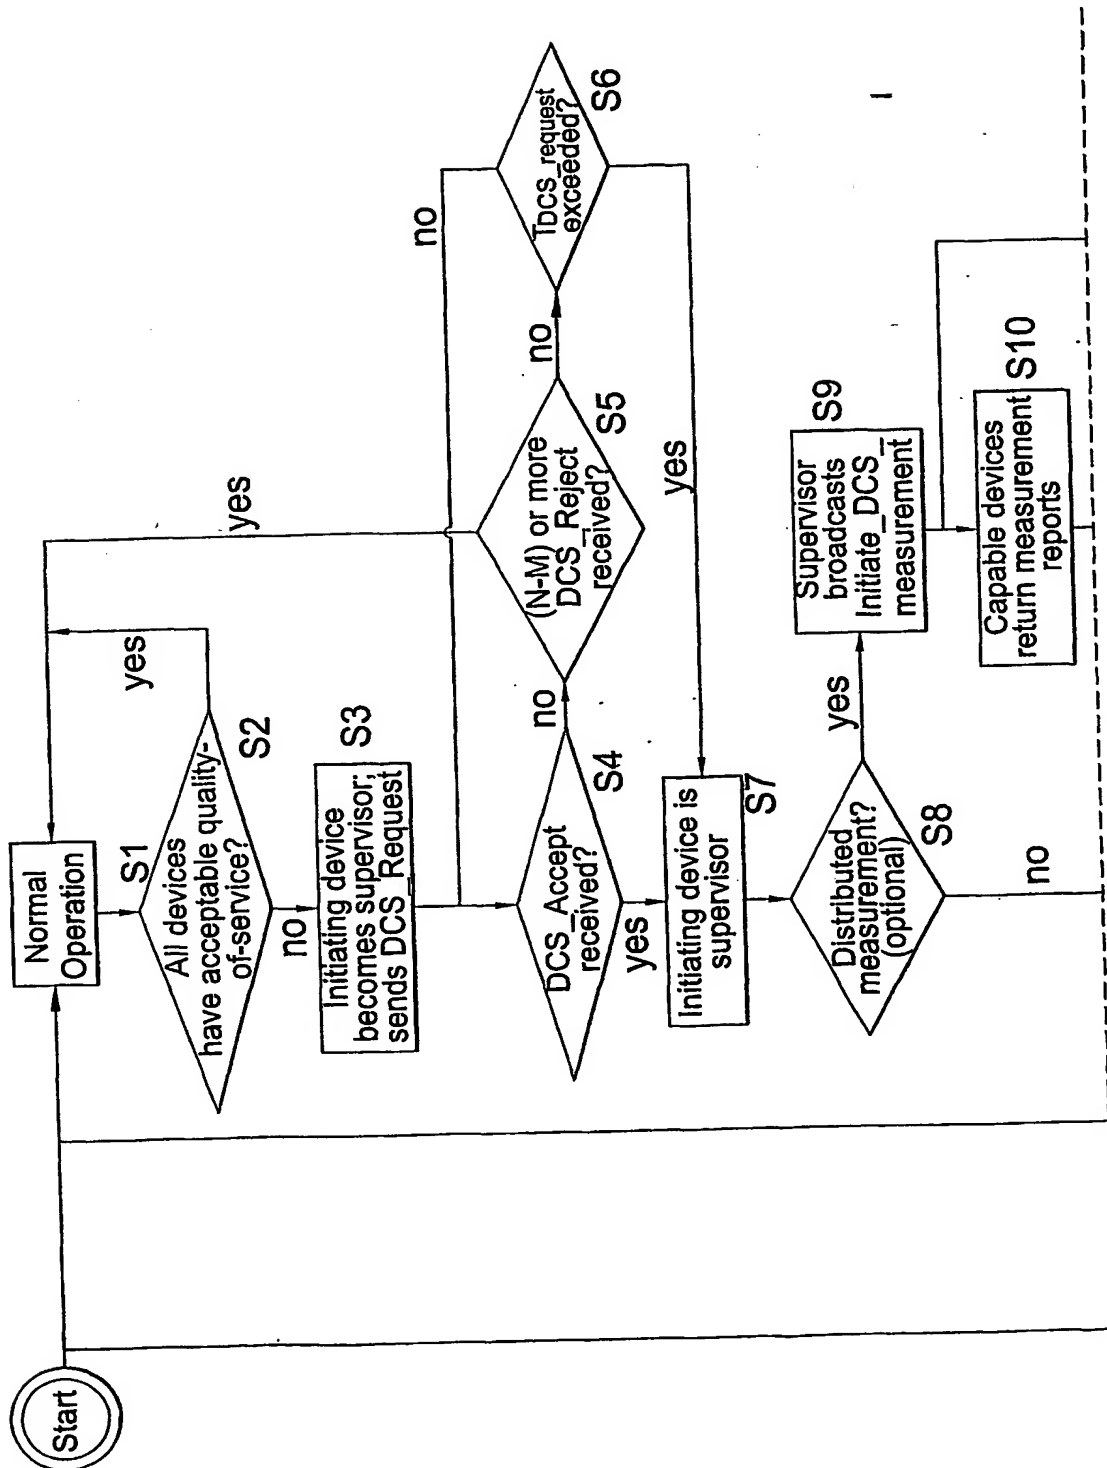


Fig. 2

3/5




Fig. 2(cont'd)

4/5

Odd number = master-to-slave slot
 Even number = slave-to-master slot




Fig. 3

5/5

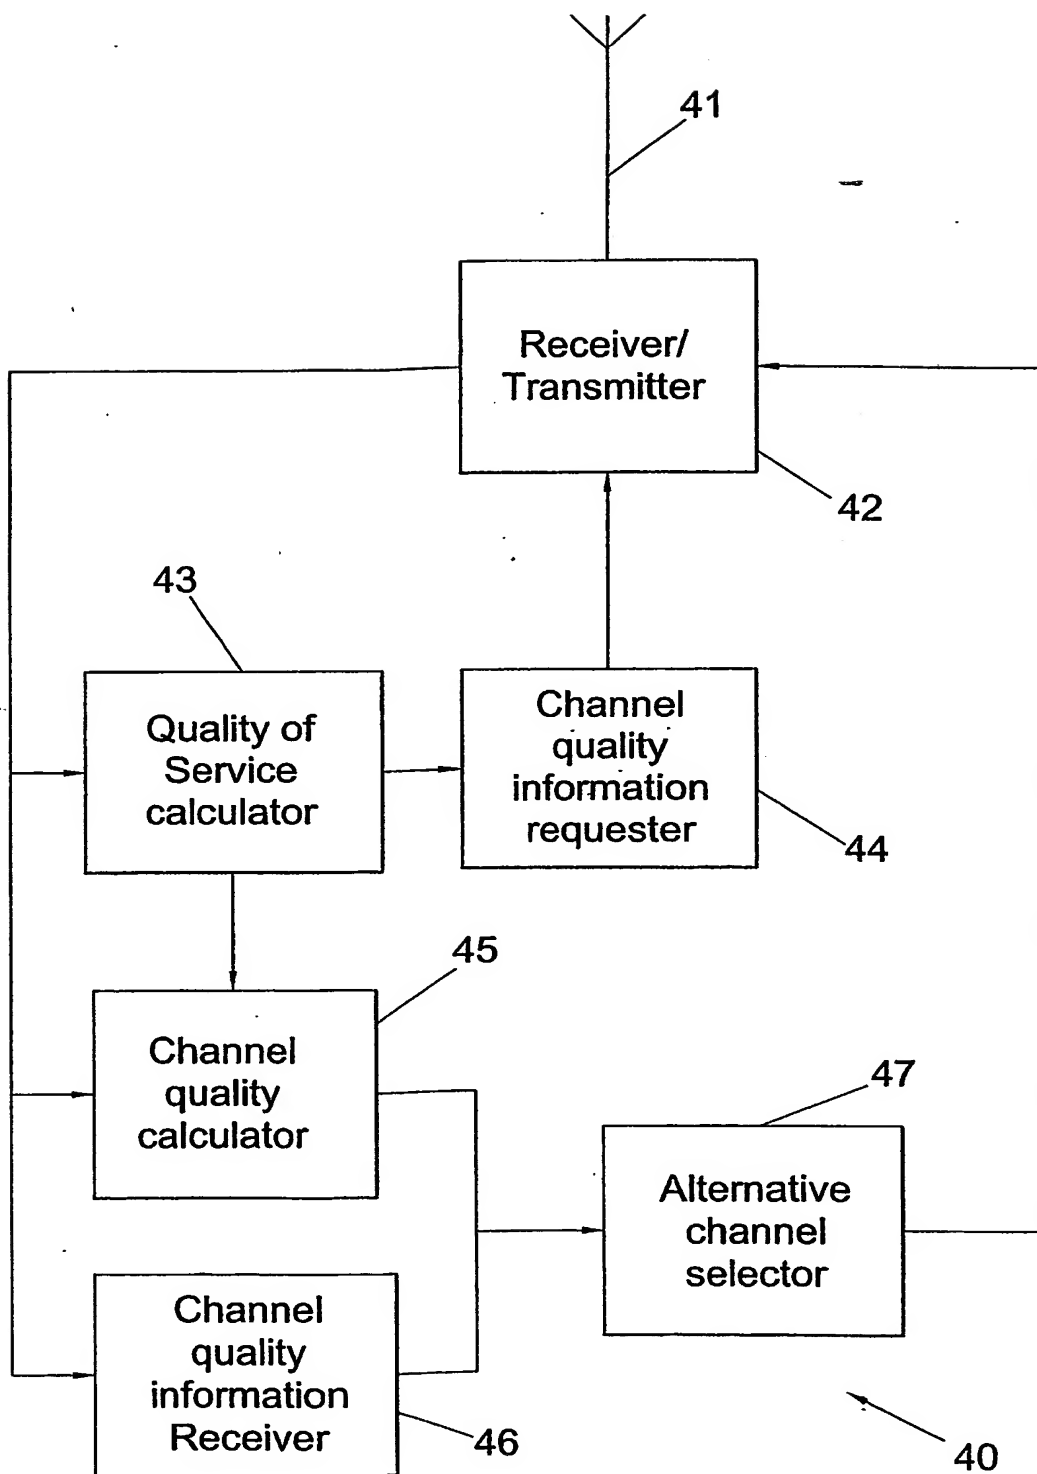


Fig. 4

(S7) Abstract: The invention relates to a system having a decentralized, distributed topology for a local area network or the like. Devices in the system monitor the channel quality and if it is determined that the quality is insufficient, a re-evaluation of the channel is carried out to try to identify a better channel. This new channel is then communicated to other devices in the network so that all devices can communicate on the new channel.



— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(88) Date of publication of the international search report:
24 April 2003

INTERNATIONAL SEARCH REPORT

Internation location No

PCT/GB 02/03929

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 079 651 A (LUCENT TECHNOLOGIES INC) 28 February 2001 (2001-02-28) paragraph '0010! - paragraph '0029!; figure 6	1-3,5-7, 13-18, 21,22
X	EP 1 119 112 A (TEXAS INSTRUMENTS INC) 25 July 2001 (2001-07-25) paragraph '0008! - paragraph '0029!; figure 1	1-3,5-7, 13,16, 17,21,22
P,X	EP 1 187 504 A (TEXAS INSTRUMENTS INC) 13 March 2002 (2002-03-13) paragraph '0005! - paragraph '0025!; figure 1 -/-	1-25

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

G document member of the same patent family

Date of the actual completion of the international search

6 February 2003

Date of mailing of the international search report

13/03/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Jaskolski, J

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/GB 02/03929

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P, X	WO 02 23819 A (KONINKL PHILIPS ELECTRONICS NV) 21 March 2002 (2002-03-21) page 5, line 1 -page 10, line 32; figures 3,4,6B -----	1-8, 10-18, 21,22

INTERNATIONAL SEARCH REPORT

Internat. application No.
PCT/GB 02/03929

Box I Observations where certain claims were found unsearchable (Continuation of item 1 of first sheet)

This International Search Report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claims Nos.: 26-28
because they relate to parts of the International Application that do not comply with the prescribed requirements to such an extent that no meaningful International Search can be carried out, specifically:
see FURTHER INFORMATION sheet PCT/ISA/210
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this International Search Report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this International Search Report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this International Search Report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International Application No. PCT/GB 02 03929

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

Continuation of Box I.2

Claims Nos.: 26-28

The wording of the claims 26-28 render it impossible to determine the matter for which protection is sought. These claims do not comply with the clarity and conciseness requirements of Article 6 PCT (see also Rule 6.1(a) PCT) to such an extent that a meaningful search is impossible. Consequently, the search has been carried out for the remaining parts of the application.

The applicant's attention is drawn to the fact that claims, or parts of claims, relating to inventions in respect of which no international search report has been established need not be the subject of an international preliminary examination (Rule 66.1(e) PCT). The applicant is advised that the EPO policy when acting as an International Preliminary Examining Authority is normally not to carry out a preliminary examination on matter which has not been searched. This is the case irrespective of whether or not the claims are amended following receipt of the search report or during any Chapter II procedure.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 02/03929

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
EP 1079651	A	28-02-2001	AU 5339800 A BR 0003576 A CN 1285696 A EP 1079651 A1 JP 2001086555 A	22-02-2001 03-04-2001 28-02-2001 28-02-2001 30-03-2001
EP 1119112	A	25-07-2001	EP 1119112 A2 JP 2001285142 A	25-07-2001 12-10-2001
EP 1187504	A	13-03-2002	EP 1187504 A2	13-03-2002
WO 0223819	A	21-03-2002	US 2002060995 A1 AU 1220202 A WO 0223819 A2	23-05-2002 26-03-2002 21-03-2002